

5

## A SYSTEM AND METHOD EMPLOYING A GRID INDEX FOR LOCATION AND PRECISION ENCODING

10

### BACKGROUND

This application claims priority under 35 U.S.C. Section 119(e)(1) of provisional application number 60/469,314, filed May 9, 2003.

15

#### Technical Field:

This invention is directed toward a system and method for location and precision encoding. More specifically, the invention is directed toward a system and method for combining the precision estimate of a database entry's coordinate value such that the precision information is included as part of a one-dimensional index.

20

#### Background Art:

25

Geographic Information Systems (GIS) are well-organized databases of information catalogued, in part, by their relationship to geography. Often, these databases contain on the order of millions or billions (or more) of data elements, each of which is associated with a particular geographical location. As an example, consider a world-wide photo database, where each photograph is associated with the location where it was taken.

30

Although modern relational databases perform a variety of algorithmic optimizations to speed up the retrieval of data contained in them, they remain inefficient at particular tasks. For example, databases are relatively poor at performing range queries over more than a single dimension. These are queries of the form, "Retrieve all data in which Value X is between  $x_1$  and  $x_2$ , Value Y is between  $y_1$  and  $y_2$ , and Value Z is between  $z_1$  and  $z_2$ ." (Again, modern databases can handle even these types of queries well, if the number of data elements is comparatively small; but as they cannot know the form of the data a priori, they necessarily are designed for generality at the expense of efficiency, and remain inefficient for large amounts of data.)

Furthermore, in some GIS applications, elements are associated with geographical regions at different scales, and queries may be performed over a range of scales. For example, in a photo database, an item may be associated with the city of New York as a whole or more specifically with the city block representing the Empire State Building. In certain kinds of queries, it may be useful to retrieve all items taken at the Empire State Building when querying for photos taken in New York City, but the converse may be undesirable – not all photos taken in New York City will be relevant to a query about the Empire State Building.

Therefore, what is needed is a system and method that can create a database of location-tagged data, for which retrieval performance is optimized for location-based queries at a particular scale and finer, no matter of the massiveness of the database.

## SUMMARY

The invention is directed toward a system and method for combining the precision estimate of a database entry's coordinate value such that the precision information is included as part of a one-dimensional index by constructing a

hierarchical index in which the size of the grid is related to the precision of the coordinate value.

5 A grid index is a gridding of an n-dimensional space into a regular partition of the grid space forming grid units, for which for a point in space,  $x$ , there is a function index ( $x$ ), which retrieves a unique integer value for the grid that contains Point  $x$ . A hierarchical grid index is effectively a number of grid indices overlaid on the same space, with grid units of different sizes. In this case, each of the functions requires an additional argument which specifies the size of the grid unit to use. Thus, assuming that the grid size,  $s$ , is drawn from a set of grid sizes,  $S$ , Index ( $x, s$ ) returns a unique integer value for the grid of size  $s$  that contains Point  $x$ . Hierarchical indexes are used to enhance the performance of database queries. A query that seeks results from a small grid size,  $s_{small}$ , does not seek matches at a large grid size,  $s_{large}$ .

15 Similar calculations can be performed for a finite area  $A$ . That is, a grid index can be used for which for a finite area in space,  $A$ , there is a function index ( $A$ ) which retrieves a unique integer value for a grid that wholly contains Area  $A$ .

20 Generally, to obtain a database wherein the precision estimate of a database entry's coordinate value is combined such that information is included as part of the one-dimensional index, one or more location entities are input into the system and method according to the invention. The grid index series is then computed wherein each grid is represented as a series of grids that incorporate the location of each location entity. This grid index series can then be output to a database.

25 More specifically, one embodiment of the system and method according to the invention uses an equirectangular projection (also known as "unprojected lat/long"), in which latitude/longitude (lat/long) values of location entities are taken as straight x-y pairs on a Euclidean coordinate system and input into the

system. The globe is then gridded at twenty different resolutions, with “square” units whose sides correspond to  $20 \times (1/2)^r$  degrees, for  $0 \leq r < 20$ . At the equator, these values correspond to scales ranging from ~240km down to ~0.5m. Next, each grid is indexed in raster-scan order. So, a given longitude/latitude coordinate  $(long, lat)$ , whose measurement error is expected to be normally distributed with standard deviation  $\sigma$  meters would be indexed as follows. First, the longitudinal span in degrees that  $3\sigma$  meters corresponds to:  $d = [180(3\sigma) \cos(lat)] / k\pi$  is determined, where  $k$  is the circumference of the earth in meters ( $4 \times 10^8$  m). Next the degree-scale of precision,  $r$ , is determined to be the discrete unit of resolution that is just larger than  $d$ :  $r = \lceil -\log_2 d / 20 \rceil$ . Finally, the coordinate  $(long, lat)$  is mapped to the index,

$$l = \left( \frac{360}{r} \right) \left\lfloor \frac{lat + 90}{r} \right\rfloor + \left\lfloor \frac{long + 180}{r} \right\rfloor$$

The coordinate as mapped to the index can then be output to a database. To recover the lat/long value, this operation is inverted:

$$lat = \frac{lr^2}{360} - 90 + \frac{r}{2},$$

$$long = l \% \frac{r^2}{360} - 180 + \frac{r}{2},$$

where ‘%’ is the modulus operator, and the  $r/2$  terms center the returned values in their grid.

The representation described above conveniently packages 2D lat/long coordinates together with precision, and it can be used to index items in a database with a single 8-byte index. If queries are issued for *all* of the items in a particular grid, it would be necessary to make multiple queries to retrieve items with location tags that are more precise than the given grid. Querying for all items taken over a large area would be computationally expensive and time consuming. To avoid this, twenty fields (one per grid resolution) are used, each of which represents the location of an item at a particular precision. For a given item with  $(long, lat)$  and precision  $\hat{r}$ ,  $l_r$  is computed, for all  $r \geq \hat{r}$ ; and, for  $r < \hat{r}$ , a

value of null is assigned. This scheme allows one to query for *all* of the items that are known to occur within a particular grid at precision  $r$ , with a single, exact-match query over the field representing location at precision  $r$ . Note that items whose location-tag precisions are coarser than that queried for will not be returned, even if the grids intersect (this inverts the standard usage of hierarchical grid indexes).

The system and method according to the invention is advantageous in that it can query very large databases in a fraction of the time it would take to generate a multidimensional query, as would be typical in traditional database schema.

In addition to the just described benefits, other advantages of the present invention will become apparent from the detailed description which follows hereinafter when taken in conjunction with the drawing figures which accompany it.

It is noted that in the remainder of this specification, the description refers to various individual publications identified by a numeric designator contained within a pair of brackets. For example, such a reference may be identified by reciting, "reference [1]" or simply "[1]". A listing of the publications corresponding to each designator can be found at the end of the Detailed Description section.

## DESCRIPTION OF THE DRAWINGS

The specific features, aspects, and advantages of the invention will become better understood with regard to the following description, appended claims, and accompanying drawings where:

FIG. 1 is a diagram depicting a general purpose computing device constituting an exemplary system for implementing the invention.

5 FIG. 2 is a simplified flow diagram of the overall process of generating a database that computes a grid index series wherein each point is represented as a series of grids according to the invention.

10 FIG. 3A is a simplified flow diagram of one embodiment of the invention wherein latitude and longitude are represented as x-y pairs on a Euclidean coordinate system and wherein the globe is gridded at 20 different resolutions.

FIG. 3B is a simplified flow diagram of the process action of indexing each grid in FIG. 3A in raster scan order.

15 FIG. 4 depicts depicts the gridding for coordinate and precision indexing employed in one embodiment of the invention. The quasi-rectangular region outlined in bold in FIG. 4(a) shows a 20-degree by 20-degree "square" on the globe. In FIG. 4(b), the same grid is projected using equirectangular projection; dark lines further subdivide into 10-degree square grids; grey lines  
20 into 5-degree grids. In FIG. 4(c) shows overlapped grids used for greater point accuracy at a given precision. Each of the larger squares represents a single grid at the same precision, but drawn from overlapped grid units. If the dots represent lat/long coordinates of items at the represented precision, then only the dots shown in bold would belong to the grid at the top left.

25 FIG. 5 depicts a flow diagram of a database query used in the system and method according to the invention.

30 FIG. 6 depicts an exemplary database structure employed in one embodiment of the system and method according to the invention.

## DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

5 In the following description of the preferred embodiments of the present invention, reference is made to the accompanying drawings that form a part hereof, and in which is shown by way of illustration specific embodiments in which the invention may be practiced. It is understood that other embodiments may be utilized and structural changes may be made without departing from the scope of the present invention.

10

### **1.0 EXEMPLARY OPERATING ENVIRONMENT**

15 FIG. 1 illustrates an example of a suitable computing system environment 100 on which the invention may be implemented. The computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computing environment 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 100.

20

25 The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

30

The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

With reference to FIG. 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a computer 110. Components of computer 110 may include, but are not limited to, a processing unit 120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The system bus 121 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus also known as Mezzanine bus.

Computer 110 typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by computer 110 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data



structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 110. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media.

10 The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as acoustic, RF, infrared and other wireless media. Combinations of

15 the any of the above should also be included within the scope of computer readable media.

The system memory 130 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 131 and random access memory (RAM) 132. A basic input/output system 133 (BIOS), containing the basic routines that help to transfer information between elements within computer 110, such as during start-up, is typically stored in ROM 131. RAM 132 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 120. By way

20 of example, and not limitation, FIG. 1 illustrates operating system 134, application programs 135, other program modules 136, and program data 137.

The computer 110 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, FIG. 1 illustrates a hard disk drive 141 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 151 that reads from or writes

30

to a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that reads from or writes to a removable, nonvolatile optical disk 156 such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

The drives and their associated computer storage media discussed above and illustrated in FIG. 1, provide storage of computer readable instructions, data structures, program modules and other data for the computer 110. In FIG. 1, for example, hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules 146, and program data 147. Note that these components can either be the same as or different from operating system 134, application programs 135, other program modules 136, and program data 137. Operating system 144, application programs 145, other program modules 146, and program data 147 are given different numbers here to illustrate that, at a minimum, they are different copies. A user may enter commands and information into the computer 110 through input devices such as a keyboard 162 and pointing device 161, commonly referred to as a mouse, trackball or touch pad. Other input devices (not shown) may include a joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus 121, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 191 or other type of display device is also connected to the system bus 121 via an interface, such as a video interface 190. In addition to

the monitor, computers may also include other peripheral output devices such as speakers 197 and printer 196, which may be connected through an output peripheral interface 195. Of particular significance to the present invention, a camera 163 (such as a digital/electronic still or video camera, or  
5 film/photographic scanner) capable of capturing a sequence of images 164 can also be included as an input device to the personal computer 110. Further, while just one camera is depicted, multiple cameras could be included as an input device to the personal computer 110. The images 164 from the one or more cameras are input into the computer 110 via an appropriate camera interface  
10 165. A microphone 168 may also be associated with the camera and input into the computer 110 with an interface 165 that is associated with one or more cameras. This interface 165 is connected to the system bus 121, thereby allowing the images to be routed to and stored in the RAM 132, or one of the other data storage devices associated with the computer 110. However, it is  
15 noted that image data can be input into the computer 110 from any of the aforementioned computer-readable media as well, without requiring the use of the camera 163.

The computer 110 may operate in a networked environment using logical  
20 connections to one or more remote computers, such as a remote computer 180. The remote computer 180 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 110, although only a memory storage device 181 has been illustrated in FIG. 1. The  
25 logical connections depicted in FIG. 1 include a local area network (LAN) 171 and a wide area network (WAN) 173, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

30 When used in a LAN networking environment, the computer 110 is connected to the LAN 171 through a network interface or adapter 170. When

used in a WAN networking environment, the computer 110 typically includes a modem 172 or other means for establishing communications over the WAN 173, such as the Internet. The modem 172, which may be internal or external, may be connected to the system bus 121 via the user input interface 160, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, FIG. 1 illustrates remote application programs 185 as residing on memory device 181. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

The exemplary operating environment having now been discussed, the remaining parts of this description section will be devoted to a description of the program modules embodying the invention.

## **2.0 A SYSTEM AND METHOD EMPLOYING A GRID INDEX FOR LOCATION AND PRECISION CODING.**

### **2.1 Overview.**

The invention is directed toward a system and method for combining the precision estimate of a database entry's coordinate value such that the precision information is included as part of the one-dimensional index by constructing a hierarchical index in which the size of the grid is related to the precision of the coordinate value. A hierarchical grid index is formed from a number of grid indices overlaid on the same space. A grid index is a gridding of an n-dimensional space into a regular partition of the grid space into grid units, for which for a point in space,  $x$ , there is a function index  $(x)$ , which retrieves a unique integer value for the grid that contains Point  $x$ . A hierarchical grid index is effectively a number of grid indices overlaid on the same space, with grid units of

different sizes. In this case, each of the functions employs an additional argument which specifies the size of the grid unit to use. Thus, assuming that the grid size,  $s$ , is drawn from a set of grid sizes,  $S$ , Index  $(x,s)$  returns a unique integer value for the grid of size  $s$  that contains Point  $x$ . Hierarchical indexes are used to enhance the performance of database queries. A query that seeks results from a small grid size,  $s_{small}$ , does not seek matches at a large grid size,  $s_{large}$ .

Similar calculations can be performed for a finite area  $A$ . That is, a grid index can be used for which for a finite area in space,  $A$ , there is a function index  $(A)$  which retrieves a unique integer value for a grid that wholly contains Area  $A$ .

FIG. 2 depicts a simplified flow diagram of the process actions of creating a database wherein the precision estimate of a database entry's coordinate value is combined such that information is included as part of the one-dimensional index. This is done by constructing a hierarchical index in which the size of the grid is related to the precision of the coordinate value. In the most general terms, the location entities are input, as shown in process action 202. The location entities can be a point or a finite area. The grid index series is then computed wherein each grid is represented as a series of grids that incorporate the location of each location entity (process action 204). This index series can then be output to a database (process action 206).

The following paragraphs more specifically discuss the system and method according to the invention.

## **2.2 Location Representation.**

Parameters other than two-dimensional location (and time) are necessary to uniquely position a real-world object. These include altitude, up to three degrees of orientation specification, and possibly other parameters, depending

on what the object is. In this discussion, examination is restricted to 2D location, as it is a manageable first step towards a broader understanding that incorporates the remaining parameters. However, it should be noted that the system and method according to the invention can process multi-dimensional location information.

*Georeferencing*, or how location is represented, is a fundamental issue in GIS. Location can be represented in many ways: established place names ("San Francisco"), user-dependent place names ("Grandma's house"), street addresses, zip codes, latitude/longitude coordinates, Euclidean coordinates with respect to some origin, and so forth.

The vast majority of Geographic Information System (GIS) projects use latitude and longitude coordinates – henceforth lat/long – with coordinates defined with respect to the WGS84 standard to specify geographic point coordinates. This scheme is followed as a way of specifying points on the globe because it offers a concise, established way to represent point location.

For the purposes of browsing and interaction, one needs two additional types of location data structures: (1) for maps and queries, a notion of a physical region that corresponds to a map as displayed, and (2) for metadata-tagging, a data type that includes both lat/long coordinates and a measure of precision or resolution.

Maps are most often displayed as rectangles on 2D displays. So, to represent a map, a structure that is called the area type is used, that is defined by a center lat/long, and width and height in kilometers, each measured from the center point and along lines of latitude and longitude. This defines a unique region on the globe that is "rectangular," circumscribed by four great-circle arcs on the 3D globe. When projected onto 2D, an area appears roughly rectangular

with sides that may deviate slightly from straight lines depending on the projection type, scale, and position on the globe.

For location tags of data items, one would like to represent lat/long and some indication of precision, to distinguish between an item associated with the Empire State Building, and an item known to be associated only with the coarser label, New York City. The priority is fast retrieval from a database over a potentially massive number of such items. In an off-the-shelf relational database, queries can be optimized if database entries can be indexed by a single number. Thus, the system and method according to the invention employs a scheme that can reduce the three continuous variables of lat/long and precision into a single, discretized index. Spatial data structures that fulfill this criterion are called *discrete global grid systems* (DGGs), of which some examples are given in Table 1.

Table 1: Candidate Spatial Data Structures for Gridding.

DGGS	Brief Description	Coordinate-Index Mapping	area-Index Mapping	Areal Variation	Shape Distortion
Lat/long gridding	Unprojected ( <i>long, lat</i> ) as cartesian x-y grid	Simple	Medium	High	High at poles
O-QTM [1]	Octahedral facets gridded by equilateral triangles	Medium	Complex	Low	Low
Dymaxion [1]	Icosahedral facets gridded by equilateral triangles	Complex	Complex	Low	Low
ISEA3H [2]	Equal-area gridding by hexagons and pentagons	Complex	Complex	None	Low

Referring now to FIGs. 3A, 3B and 4, one embodiment of the system and method according to the invention uses an equirectangular projection (also

known as “unprojected lat/long”), in which lat/long values are taken as straight x-y pairs on a Euclidean coordinate system (see FIG. 4(b)) and are input into the system (FIG. 3A, process action 302). As shown in process action 304, the globe is gridded at twenty different resolutions, with “square” units whose sides correspond to  $20 \times (1/2)^r$  degrees, for  $0 \leq r < 20$ . FIG. 4(b) shows 20-, 10-, and 5-degree gridding of the region outlined in FIG. 4(a). At the equator, these values correspond to scales ranging from ~240km down to ~0.5m. The gridding can occur at greater or fewer resolutions, if more or less precision is required.

Next, as shown in process action 306, each grid is indexed in raster-scan order. So, a given longitude/latitude coordinate (*long*, *lat*), whose measurement error is expected to be normally distributed with standard deviation  $\sigma$  meters would be indexed as follows. First, as shown in FIG. 5, process action 306A, the longitudinal span in degrees that  $3\sigma$  meters corresponds to:  $d = [180(3\sigma) \cos(lat)] / k\pi$  is determined, where  $k$  is the circumference of the earth in meters ( $4 \times 10^8$ m). Next the degree-scale of precision,  $r$ , is determined to be the discrete unit of resolution that is just larger than  $d$ :  $r = \lceil -\log_2 d \rceil$ , as shown in process action 306B. Finally, as shown in process action 308, the coordinate (*long*, *lat*) is mapped to the index,

$$I = \left( \frac{360}{r} \right) \left\lfloor \frac{lat + 90}{r} \right\rfloor + \left\lfloor \frac{long + 180}{r} \right\rfloor \quad (1)$$

The coordinate as mapped to the index can then be output to a database (process action 310). To recover the lat/long value, this operation is inverted:

$$\begin{aligned} lat &= \frac{Ir^2}{360} - 90 + \frac{r}{2}, \\ long &= I \% \frac{r^2}{360} - 180 + \frac{r}{2}, \end{aligned} \quad (2)$$

Where ‘%’ is the modulus operator, and the  $r/2$  terms center the returned values in their grid. Of course, because of the floor operations in Eq. (1), one can only recover the value to  $\pm 1/2$ , which conveniently is the precision of the item’s location estimate. As  $r$  decreases, the precision asymptotically approaches zero,



as desired. It is noted that these grids do *not* generally correspond to physical area objects, but the grid units that occur in a particular area can be easily determining by enumerating units that overlap with an area. If greater accuracy is needed, overlapping grids at each scale could be used, with coordinates mapped to the square whose center is closest (see FIG. 3(c)).

### 2.3 Grid Types.

Using a square grid with equirectangular projection is intuitive and keeps database queries efficient. But, it is inelegant in that the physical size and shape of grid units at the same indexing scale is distorted, particularly near the poles. Other gridding schemes can be used, however. In one embodiment of the invention a gridding scheme is employed that has minimal areal variation and shape distortion and uses a hierarchical series of equilateral polygons embedded within a Platonic solid. The major advantage of this gridding scheme is that at a fixed scale, units are of similar size and shape. Naturally, the equations for determining the grid index given a coordinate and a coordinate given a grid index must be adapted for these alternate schemes.

### 2.3 Location Database Schema Design

The representation described above conveniently packages 2D lat/long coordinates together with precision, and it can be used to index items in a database with a single 8-byte index.

If queries are issued for *all* of the items in a particular grid, it would be necessary to make multiple queries to retrieve items with location tags that are more precise than the given grid. Querying for all items taken over a large area would be computationally expensive and time-consuming.

To avoid this, twenty fields (one per grid resolution) are used, each of which in one embodiment represents the location of an item (e.g., location entity) at a particular precision. For a given item at (*long*, *lat*) and precision  $r^*$ ,  $l_r$  is computed as in Eq. 1, for all  $r \geq r^*$ ; and, for  $r < r^*$ , a value of null is assigned.

5 This scheme allows one to query for *all* of the items that are known to occur within a particular grid at precision  $r$ , with a single, exact-match query over the field representing location at precision  $r$ . Note that items whose location-tag precisions are coarser than that queried for will not be returned, even if the grids intersect (this inverts the standard usage of hierarchical grid indexes). That is, if  
10 one is searching for all items taken within a certain Manhattan block, a query for that grid unit will not return an item about which is known only that it was associated somewhere in New York City.

A flow diagram of the process actions of retrieving the desired data from a  
15 database of one database configured by the system and method of the invention is shown in FIG. 5. A user queries the database to return all location entities in a given cell  $y$  at a given scale  $S$ , as shown in process action 502. Column  $S$  of the database corresponding to the desired scale is then searched for values of cell  $y$  (process action 504) and these values are returned (process action 506). Note  
20 that *all* entities in the database table which occur within the grid represented by cell  $y$  will be included in the retrieval results, even those which occur at scales finer than cell  $y$  itself.

An exemplary database format for one embodiment of the invention is  
25 shown in FIG. 6. The database contains a location entity ID 602 for each entry in the database. Optionally, the latitude and longitude coordinates 604, as well as an indicator representing the grid scale 606, can also be recorded. The scale index at scale 1 provides a single integer that is the grid at a scale 1 that wholly contains the location representation or is set to null if it is not contained (as  
30 would be the case if the location representation were a finite area or region instead of a finite point). The scale index at scale 2 provides a single integer that

represents a grid cell at a finer resolution than scale 1. And, so forth until scale N. If the finite area is not wholly contained within a given grid cell the scale index is set to a null value for this location representation.

- 5            Although this representation requires additional fields in the database, this is outweighed by the gain in performance (nearly tenfold for >1 million rows).

10           The foregoing description of the invention has been presented for the purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching. It is intended that the scope of the invention be limited not by this detailed description, but rather by the claims appended hereto.

## References:

- [1] Diomidis, D. S., Position-annotated photographs: a geotemporal web. *IEEE Pervasive Computing*. 2(2):72-79, 2003.
- [2] Sahr, K. and White, D., Discrete Global Grid Systems. In *Proc. 13th Symp. Interface, Comp. Sci. & Stat.* 269-78, 1998.

5